Attachment 17 to Section J
ERA Deployment Concept

DRAFT

1.0 Summary

The ERA system, because of its size, scope, and commitment to long term preservation and servicing of government records, will require deployment and design approaches that support its unique nature and mission goals. When designing and deploying ERA, NARA must take a long term view for the system's operation and its required scalability, reliability, and cost effective operations. This long term vision will accommodate the use of outsourcing of potential processing and hosting services while at the same time ensures NARA's stewardship of the records entrusted to it. This document outlines concepts behind the development and deployment of ERA. This document is not meant to prescribe a design, number of required facilities, or a rigid architecture. The concepts and design drivers outlined in the document are what should be noted. Highlights of the proposed ERA deployment strategy include:

1. NARA will contract for a core ERA system – The ERA core system is a group of instances that reliably preserve and service a full set of all electronic records holdings for which NARA is responsible. The core system includes any classified instances that may be required

2. The designated installation sites are controlled, but not necessarily owned, by the government

3. NARA will, at its discretion, make the core system software and specifications available to industry, academia, and other government agencies (in accordance with appropriate government software and data rights). This will support:
   a. Hosting of ERA instances to meet peak access loads
   b. Support of ingest or record transformation tasks
   c. Hosting of sensitive or classified records in a secure environment
   d. Allow for value added services to be developed and offered on record collections by a third party, either public or private

2.0 Statement of Deployment Goals

There are certain assumptions and drivers that sculpt the deployment approach for ERA. These assumptions and design drivers are collectively considered the deployment goals for the program. These goals include:

- NARA must own and control at least one set of all holdings of electronic records entrusted to it. This is required for protection of the records and fulfillment of NARA's mission to ensure long term preservation and access to the government's records

- The ERA system is one of NARA's contributions to the Federal Enterprise Architecture (FEA) and fulfills a critical role in the development and deployment of NARA's own Enterprise Architecture (EA)

- The design and deployment vision for ERA must allow for the contracting out of record processing and access support, if NARA chooses to exercise that option in the future. The contracting out of record services must be done within the context of NARA's mission and ultimate responsibility for the integrity of the records

- Minimize government ownership of equipment and facilities. This desire must be balanced against NARA's stewardship of the records and commitment to FEA support

- Allow industry and academia to provide value added services on record holdings

- Produce a highly reliable system design. Characteristics of such a design include:
    - Avoidance of single point/site of failure
    - Graceful performance degradation of the system when failures occur
    - Maintain system operations in face of Remedial Maintenance (RM), Preventative Maintenance (PM), and planned upgrades/changes

These deployment goals must be kept in mind when considering design approaches for ERA. A system design approach for ERA could either hinder or easily allow for flexible deployment options as ERA matures and is fielded. The rest of this paper outlines design concepts and how they can be applied to present flexible deployment options to NARA.

3.0 Architecture/Domain Drivers

In addition to the deployment goals, the design and deployment of ERA must take into account certain architectural demands and aspects of the ERA record preservation domain itself. These drivers must be accommodated in any deployment and design strategy for the ERA system. These drivers include:

- The size of the ERA record holdings. Analyses detailed in the ERA Load Analysis Report (LAR) project the ERA permanent records holdings to be in excess of 100 PBs of data 12-15 years after deployment, with continued growth in holdings in subsequent years. The sheer volume of data that must be accommodated, as well as its associated access loads, is a huge driver that must be accounted for in any architectural approach proposed. Architectural concepts including distributed deployment(s), load balancing techniques, and multiple sources for access to high demand records are applicable to the holdings size aspect of ERA.

- Ensuring the integrity of the record holdings. The records must be protected from loss, alteration, or the lack of access capability over time. Appropriate security and accommodation of timely backup of holdings with subsequent restoration of access are techniques that are required in this area.

- The evolutionary nature of the ERA system. This aspect is most pronounced in two areas:
    - Changes to the Persistent Preservation approaches used for records. Over time electronic records will need to be stored, represented, and accessed in different ways given the forward march of computer technology and the rapid obsolescence of formats and techniques.
    - Independent of the record preservation techniques, the general infrastructure and support technologies used in ERA will need to be updated and upgraded over time. Technology insertion into the ERA design will be imperative.

- The heterogeneity of holdings ERA will complicate storing and providing access to the holdings, and preserving them. The scope of this issue can be appreciated by considering that ERA records can be classified via three different attributes.
    - Record Types (RTs) – Any record will be classified according to its intellectual format. Examples of record types include letters, ledgers, maps, reports, etc.
    - Data types (DTs) – A data type is a set of lexical representations for a corresponding set of values. The values might be alphabetic characters, numbers, colors, shades of gray, sounds, et al. The lexical representation of such values in digital form assigns each value to a corresponding binary number, or string of bits. A data type may be simple, such as the ASCII representation of alphabetic characters, or composite; that is, consisting of a combination of other data types. An electronic record consists of one or more digital components; i.e., strings of bits that each has a specific data type.
    - Varied classes/collections of holdings – Records of the same RT and DT may still belong to different record series or collections, which further define the nature of the record. Examples of high level collections or series could include particular Presidential

collections, Federal record series, and potentially record collections in Federal Record Centers (FRCs).

4.0 Architectural Characteristics/Concepts

In order to begin to address the constraints and drivers outlined above, the following architectural concepts are proposed for ERA. Taken as a whole, these architectural characteristics lay the groundwork for a flexible, evolvable design for ERA.

- ERA is predicated on a Service Based Architecture – The ERA system will be built to provide services on records. These services define the extent and ability to access, manipulate, and process the records. Hardware and software COTS and developed components will interact to provide the ERA services. In fact, the ERA system can be viewed as basically the summation of the data held and the services offered on the data. Examples of services are provided in the discussion of the next architectural concept.

- Adoption of a component based design approach - Components offer the potential to assemble applications more rapidly. A key to assembling applications quickly is the ability to reuse existing pre-built components to meet your application requirements. For the purpose of this discussion, there are two key concepts for which it is necessary to have a common understanding:

    - A component is an implementation unit. It is the deployment of one or more interfaces.
    - An interface is how a consumer of a component views that component. For the purpose of consuming a component, the consumer is concerned with the interface it is consuming. Since the component is the implementation of an interface, the consumer is really concerned with the way the interface behaves.

  The interfaces associated with a component are the services provided by that component. Systems have been built around the concept of component-based design for years. Using a component based approach for design allows for:

    - Designing loosely coupled system components that encapsulate behaviors of the component
    - A rigorous definition of Application Programming Interfaces (APIs)
    - Enforcing information hiding within design components

  To illustrate the component/service based concepts, consider the OAIS Reference Model for archival systems, shown in Figure 1. Note that the following example is used for illustrative purposes only. Also, the example given addresses only the highest level services for the illustrative system design. Components called out in the example would in fact be comprised of lower level components in a true detailed design.

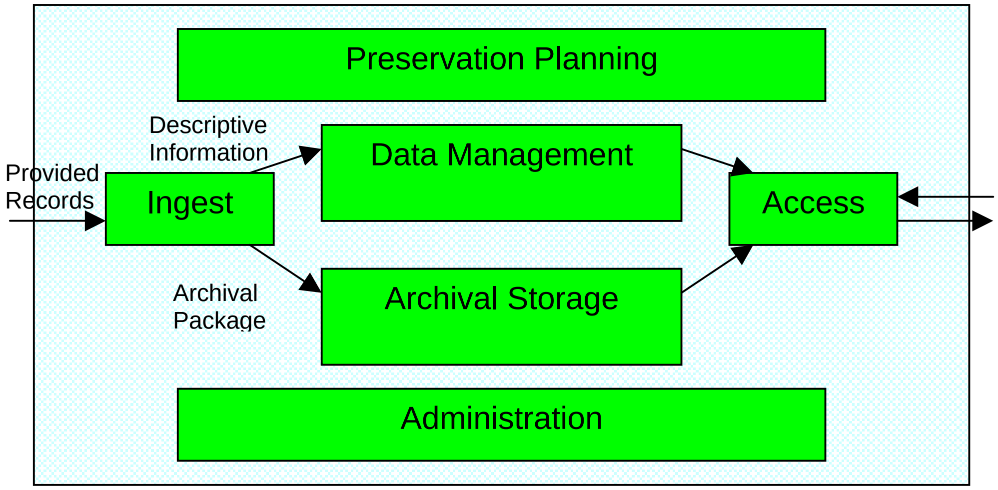

Figure 1: OAIS Reference Model for Archival Systems

Consider that the Ingest functionality called for in the OAIS model, for the purposes of this example, will be provided by a high-level component known as Ingest. The Ingest component needs to perform a number of functions within the system, the primary system functions being the acceptance of information/records provided by a producer, production of Descriptive Information, and the production of a record package ready for archival. *How* the Ingest component performs its processing of input records to produce the Descriptive Information and Archival Package is hidden from the rest of the system. The system simply sees the service interfaces provided by the Ingest component. The Ingest component needs to support a number of services, available both externally (public) and internally (private), to perform its role in the system. These services could include processing functions such as:

- Validate Transfer
- Validate Records Contents
- Transfer records from media (public)
- Transfer records electronically (public)
- Classify records
- Persistently Preserve Record
- Encapsulate Original Record
- Produce Descriptive Information (public)
- Produce Archival Package (public)

The design of the Ingest component, for this example, must support each and every public service defined above in order to function in the system. These public services, as well as perhaps some select private services, would constitute the definition of ERA *core services*. Core services are those services that a component must provide, in the prescribed manner, to be considered an ERA compliant component. The ERA core services would be the basis for NARA's contribution to the FEA.

How a component supports its public and private services is a design decision, and can in fact change over time. This is how technology insertion can be supported within the design. The overall ERA architecture and services remain consistent while the implementation can change over time using better, more efficient technologies.

- ERA services are capable of being distributed and are relocatable – A system designed for a service based approach requires some way for a component's services to be made known to the system and be invoked by other components. This need is fulfilled by the use of middleware. Middleware supports the discovery and binding of services between components. The middleware 'brokers' services between the components and supports the passing of control and data between the components during service invocation. A powerful characteristic of a brokered architecture is the ability to physically locate or relocate components and their services away from other components, be that in simply different machines in the same physical location or in a geographically dispersed manner. The ability to relocate services supports the goals of outsourcing the hosting of ERA services and the contracting for value added services to be provided by third parties. Value added service providers would produce components that would make available additional services outside the ERA core services.

- Adopt the use of an Active Safe Store approach – ERA must protect its records holdings against catastrophic failure of a portion of the system, a system malfunction/virus that corrupts the records, malicious activities that could harm the records, or potentially the compromise/destruction of an entire installation site. The classic approach to protect data against these types of threats is to maintain an off site safe-store copy of all the holdings that can be tapped for recovery if required. Possible approaches to supporting an off site copy of records can vary in complexity and cost. A few considerations must be kept in mind when considering a safe store approach.

    - Recovery time and ease - The size of the ERA holdings make full recovery of an archive's holdings expensive and time consuming. Consider the challenge of extracting 50-100 PBs from an off site data store. The time spent to read all the data out would be an obstacle in itself, preceded by the time it would take to build a new archive system to hold the records. The amount of time it would take to re-establish service is considerable.

- Maintenance of the off site holdings – The ERA system recognizes the issue of technology obsolescence, both in terms of computer resources and data formats. Any system hosting a safe store copy of data would need to accommodate technology refresh and media/records management, just as primary records storage does. (The methods used to accommodate technology refresh and media/records management may or may not be identical to the access system's approaches.) This includes computer technology used and media migration and maintenance at a minimum. From a consistency point of view, the safe store system may wish to stay in lock step with the operational record storage system technology wise.
- Synchronization with the operational record stores – The safe store system's records need to be maintained in a synchronized fashion with the primary record storage system. As records are added, deleted, or persistently preserved, the safe store system must be kept in alignment with the state of the operational records. This requires a fairly high degree of communication and data movement between the operational and safe store systems. (Note that the acceptable "lag" between primary and safe store site updates, as well as the granularity of these updates, is a cost/benefit decision requiring detailed analysis.)
- The frequency with which holdings change . - ERA will maintain one large set of holdings which, in principle, should never change; i.e., the files of permanent electronic records in the formats in which they were received from originators. ERA will also store another class of holdings which, while smaller than the first, at least in number of objects, will rarely change; i.e., the files of permanent electronic records which are in persistent archival formats.

One architectural approach that addresses these concerns is "active safe store." Active Safe Store is defined as the system's ability to actively manage the safe store copy of the data, as opposed to the passive storage of media in a vaulting or "salt mine" approach. The use of Active Safe Store implies that the system as a whole maintains the technology and record data in synch with the operational copy of the records. ERA could build dedicated Active safe store sites, however an alternative to implementing an Active Safe Store approach in a cost effective manner is to designate individual ERA sites to act as the Active Safe Store for other ERA sites. Figure 2 illustrates this approach using a three node model.

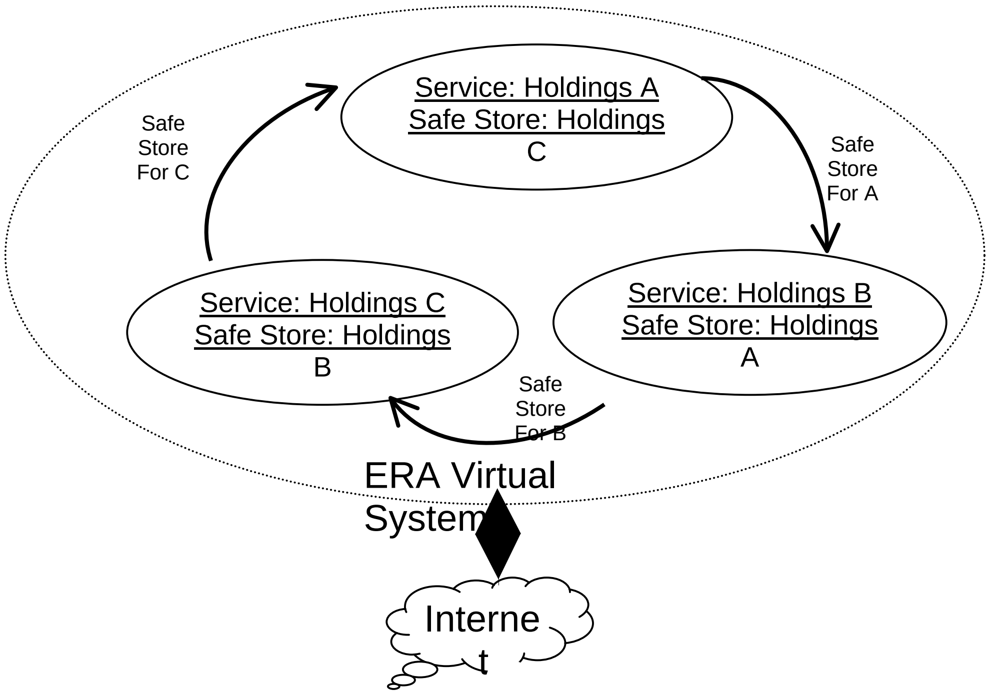

Figure 2: Active Safe Store Approach using Three Nodes

- System Availability - An adequate degree of hardware and software redundancy will be provided in order to achieve reliable services. System availability is viewed as the availability of services, however they are supported. Performance degradation for services is acceptable under failure conditions of hardware and/or software. Failure conditions may result in some services or record types becoming temporarily unavailable or impacted. It is the total disappearance of system services against all records that must be guarded against.

  The following example illustrates this concept. Assume a record search service is provided by a database. A system could be constructed where a parallel database engine was running on a cluster of hardware nodes supported by a Storage Area Network (SAN) for disk storage. If a server failed in this configuration, the parallel database would reconfigure itself and redistribute load across the remaining server(s) without appreciable service interruption. The SAN itself could be configured as a high reliability resource, where a single disk/controller failure is functionally transparent to operations (although disk rebuilds would impact performance until recovery is complete). As illustrated, single failures in the hardware configuration do not result in complete loss of functionality. Lose enough clustered servers or disks and you will of course suffer lose of functionality, but this is a cost/benefit consideration of a proposed design.

5.0 ERA: A System of Instances

The ERA system will be built up from components that provide ERA services. When appropriate and sufficient components are integrated with each other and make their services available, an *instance* of the ERA system exists. An instance of ERA is defined as a configuration of components that provides ERA services against specific record holdings. These components consist of hardware, COTS software, configuration data, and application software. An example of an ERA instance would be the integrated components that provide ERA services against Census data. The ERA Census instance would include all ERA services necessary to ingest, preserve, maintain, and provide access to the Census holdings under its control. These record holdings would encompass the actual records themselves, all Life Cycle data associated with them, and historical logs.

If an ERA instance provided all ERA core services against its record holdings, in other words if the instance was capable of servicing the entire life cycle for the records, then the instance would be considered a *full instance*. The Census example given above would be a full instance. If components were integrated into a system that provided less than the full suite of ERA record lifecycle services, then this instance would be considered a *limited instance*. An example of a limited instance would be a system whose components just provided record access services to holdings that were supplied. Continuing to use Census records as an example, NARA could designate such a limited system as a "Census access instance."

Just as components make up system instances, the overall ERA system is a number of interacting and independent instances. The ERA system model is a number of interacting components that, taken as a whole, reliably provide a full compliment of services against all record holdings. Referring back to Figure 2, each of the holdings for the systems can be viewed as a record set, with each system oval representing an ERA instance that services that set. Sets could include things like DOMPF records, Census data, State Department Cables, or Presidential materials.

6.0 ERA Deployment Approach

The architectural concepts developed in this paper provide the background to understand the proposed ERA development and deployment approach.

1) NARA will contract for the core ERA system – The ERA core system is a group of instances that, in the aggregate, provide all services required for the lifecycle management of records and reliably preserve and service a full copy of all electronic records for which NARA is responsible. The core system includes any classified instances that may be required.
    a. The ERA core system Development contractor will:
        1. Develop overall ERA architecture
        2. Design the core system
        3. Procure and install all hardware required for core system
        4. Develop software necessary to provide ERA core services
        5. Integrate COTS software, developed application software, and hardware into a functional core system
        6. Deploy instances of ERA at designated installation sites to provide for an overall system solution providing services against all required records
    b. The Development contractor for the ERA core system will provide services for Operation and Support of the system to:
        1. Operate the instances of the core ERA system at the fielded sites
        2. Provide for hardware PM and RM for the core system
        3. Coordinate with the development contractor with regard to software or hardware upgrades, installations, or changes

2) The designated installation sites are controlled, but not necessarily owned, by the government. Candidate types of installation sites include:
    a. Facilities leased by NARA
    b. Facilities owned by a different government
    c. NARA owned facilities agency

3) NARA will, at its discretion, make the core system software and specifications available to industry, academia, and other government agencies (in accordance with appropriate government software and data rights). This will allow NARA to minimize the amount of hardware, infrastructure, and facility resources it must purchase outright. Outsourcing of ERA work might include:
    a. Hosting of ERA instances to meet peak access loads. Example: When a Decennial Census becomes available there is usually a short term heavy interest in the data. This access load decreases over time. In order to support the temporarily heavy access load NARA may wish to contract out for hosting of ERA Census access instances.
    b. Support of ingest or record transformation tasks. If NARA requires a record collection to be transformed or in some way preserved, it may choose to contract the work out. Again, limited ERA instances could be hosted and utilized by a contractor to perform these tasks.

c. Hosting of sensitive or classified records in an existing secure environment. This use of this environment, provided by a different government agency or private provider, would relieve NARA of the burden of establishing and accrediting secure facilities of its own. In this arrangement a full ERA instance capable of servicing the classified/sensitive records would be hosted at the secure site.
d. Allowing for value added services to be developed and offered on record collections by a third party, either public or private. An example would be where an academic institution develops a unique search or presentation algorithm to be used against certain record types or collections. An ERA instance could be deployed with the desired records with the new algorithms integrated into the component framework. These algorithms could conceivably make their way, in a controlled and managed fashion, into the core ERA system if NARA so chose.

Figure 3 illustrates a representative deployment approach proposed for ERA. This approach uses the concepts of Active Safe Store, a multi-node core system that hosts ERA instances, outsourced ERA instances, and secured ERA core instances for classified information. Each of the core installation sites would contain one or more complete instances of ERA. The non-core (contracted) sites would hold one (or perhaps more) instances of ERA. The contracted instances of ERA shown in Figure 3 are designated as limited Access instances; however Ingest, transformation, or full ERA instances are not precluded from being contracted out. The core instances however would always be under government control.

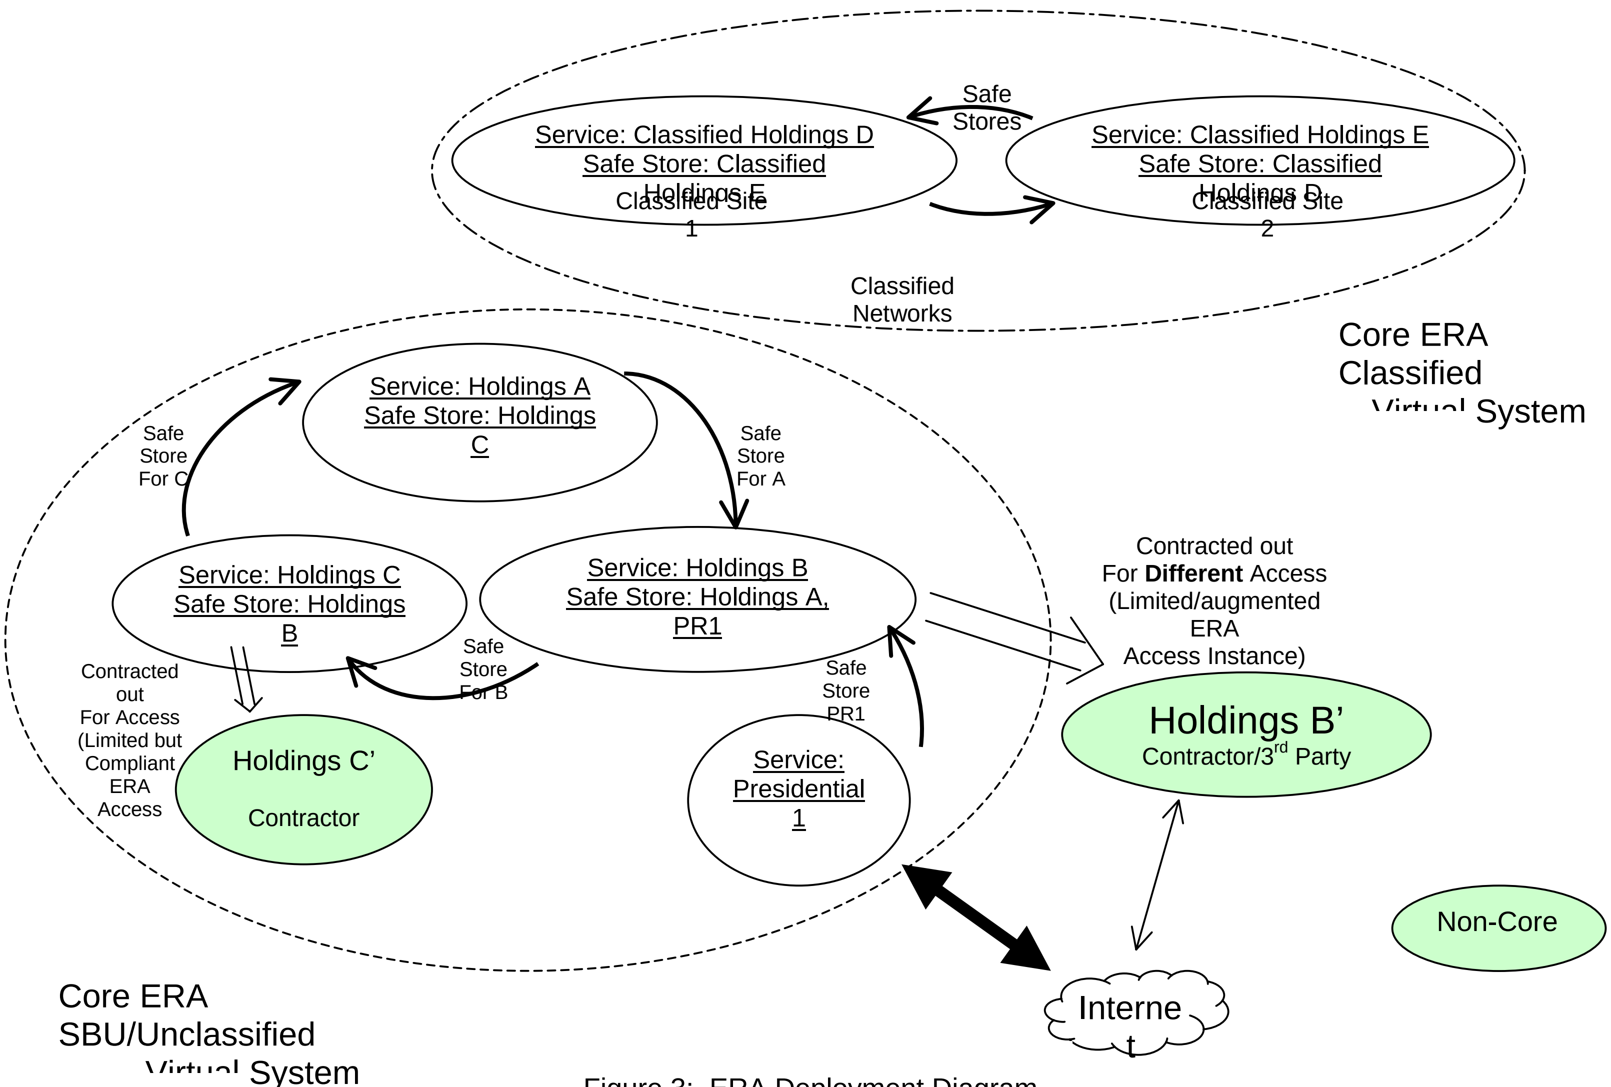

Figure 3: ERA Deployment Diagram